Cert4Prep
Practice Test Software

Amazon

MLA-C01

AWS Certified Machine Learning Engineer - Associate

Version: 10.2

Questions & Answers DEMO PDF

(Preview content before you buy)

Check the full version using the link below

https://cert4prep.com/exam/mla-c01

Unlock Full Features

Stay Updated: 90 days of free exam updates
Zero Risk: 30-day money-back policy
Instant Access: Download right after purchase
Always Here: 24/7 customer support team

Question 1. (Single Select)

Case Study

A company is building a web-based AI application by using Amazon SageMaker. The application will provide the following capabilities and features: ML experimentation, training, acentral model registry, model deployment, and model monitoring.

The application must ensure secure and isolated use of training data during the ML lifecycle. The training data is stored in Amazon S3.

The company needs to use the central model registry to manage different versions of models in the application.

Which action will meet this requirement with the LEAST operational overhead?

A: Create a separate Amazon Elastic Container Registry (Amazon ECR) repository for each model.
B: Use Amazon Elastic Container Registry (Amazon ECR) and unique tags for each model version.
C: Use the SageMaker Model Registry and model groups to catalog the models.
D: Use the SageMaker Model Registry and unique tags for each model version.

Answer: C

Explanation:
Amazon SageMaker Model Registry is a feature designed to manage machine learning (ML) models throughout their lifecycle. It allows users to catalog, version, and deploy models systematically, ensuring efficient model governance and management.
Key Features of SageMaker Model Registry:
Centralized Cataloging: Organizes models into Model Groups, each containing multiple versions.
Metadata Association: Attach metadata such as training metrics and performance evaluations to models.
Approval Status Management: Allows setting statuses like PendingManualApproval or Approved to ensure only vetted models are deployed.
Seamless Deployment: Direct integration with SageMaker deployment capabilities for real-time inference or batch processing.
Implementation Steps:
Create a Model Group: Organize related models into groups to simplify management and versioning.
Register Model Versions: Each model iteration is registered as a version within a specific Model Group.

Set Approval Status: Assign approval statuses to models before deploying them to ensure quality control.
Deploy the Model: Use SageMaker endpoints for deployment once the model is approved.
Benefits:
Centralized Management: Provides a unified platform to manage models efficiently.
Streamlined Deployment: Facilitates smooth transitions from development to production.
Governance and Compliance: Supports metadata association and approval processes.
By leveraging the SageMaker Model Registry, the company can ensure organized management of models, version control, and efficient deployment workflows with minimal operational overhead.
AWS Documentation: SageMaker Model Registry
AWS Blog: Model Registry Features and Usage

Question 2. (HOTSPOT)

A company stores historical data in .csv files in Amazon S3. Only some of the rows and columns in the .csv files are populated. The columns are not labeled. An ML

engineer needs to prepare and store the data so that the company can use the data to train ML models.

Select and order the correct steps from the following list to perform this task. Each step should be selected one time or not at all. (Select and order three.)

• Create an Amazon SageMaker batch transform job for data cleaning and feature engineering.

• Store the resulting data back in Amazon S3.

• Use Amazon Athena to infer the schemas and available columns.

• Use AWS Glue crawlers to infer the schemas and available columns.

• Use AWS Glue DataBrew for data cleaning and feature engineering.

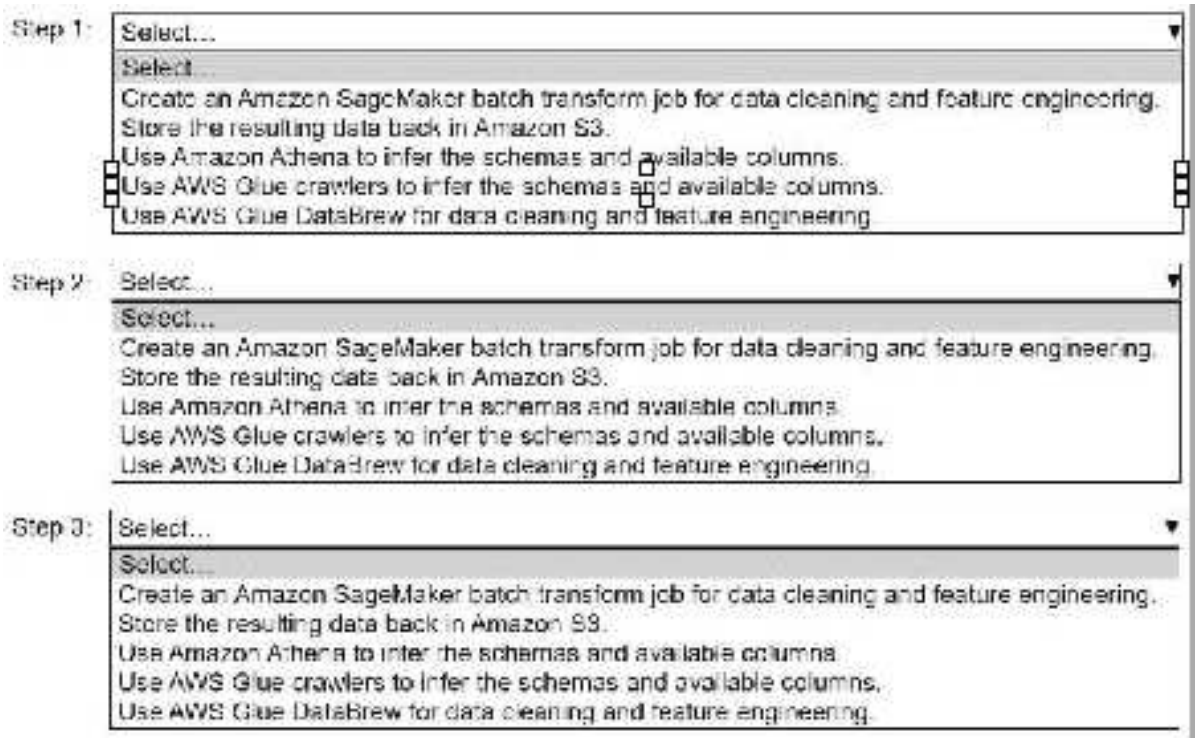
Step 1:
Select...
Select...
Create an Amazon SageMaker batch transform job for data cleaning and feature engineering.
Store the resulting data back in Amazon S3.
Use Amazon Athena to infer the schemas and available columns.
Use AWS Glue crawlers to infer the schemas and available columns.
Use AWS Glue DataBrew for data cleaning and feature engineering
Step 2:
Select...
Select...
Create an Amazon SageMaker batch transform job for data cleaning and feature engineering.
Store the resulting data back in Amazon S3.
Use Amazon Athena to infer the schemas and available columns.
Use AWS Glue crawlers to infer the schemas and available columns.
Use AWS Glue DataBrew for data cleaning and feature engineering.
Step 3:
Select...
Select...
Create an Amazon SageMaker batch transform job for data cleaning and feature engineering.
Store the resulting data back in Amazon S3.
Use Amazon Athena to infer the schemas and available columns.
Use AWS Glue crawlers to infer the schemas and available columns.
Use AWS Glue DataBrew for data cleaning and feature engineering.

Answer:

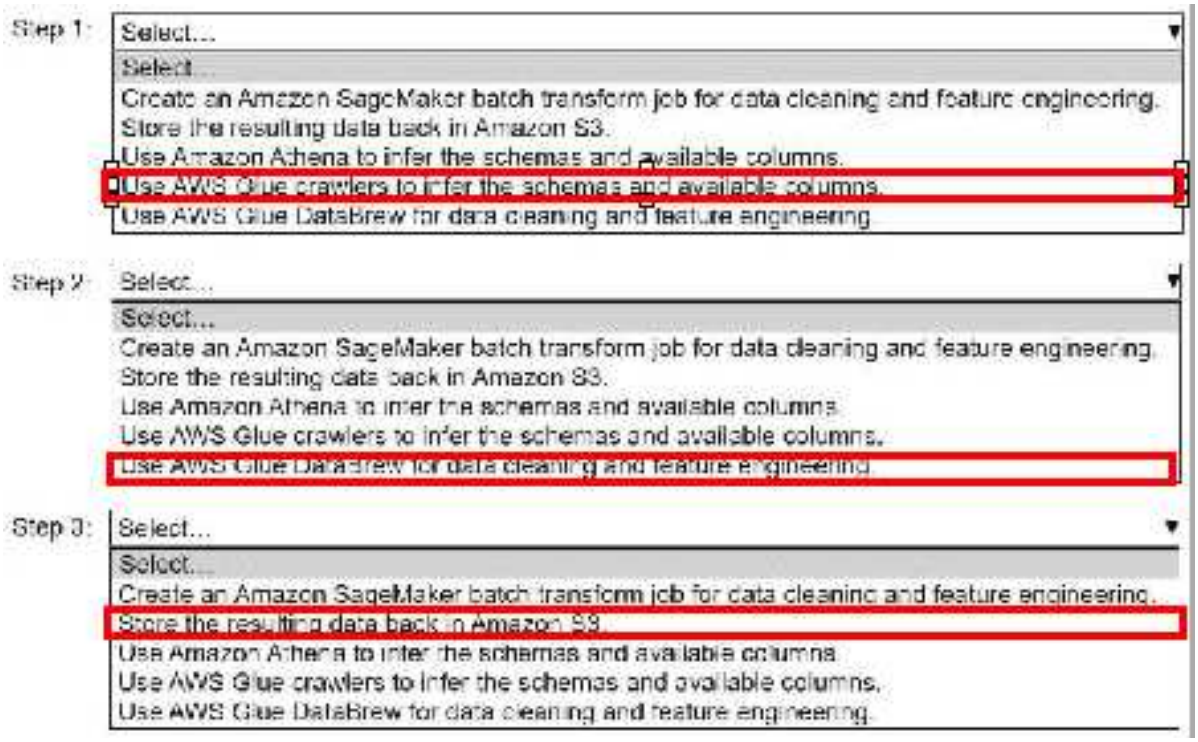
Step 1:
Select...
Select...
Create an Amazon SageMaker batch transform job for data cleaning and feature engineering.
Store the resulting data back in Amazon S3.
Use Amazon Athena to infer the schemas and available columns.
Use AWS Glue crawlers to infer the schemas and available columns.
Use AWS Glue DataBrew for data cleaning and feature engineering
Step 2:
Select...
Select...
Create an Amazon SageMaker batch transform job for data cleaning and feature engineering.
Store the resulting data back in Amazon S3.
Use Amazon Athena to infer the schemas and available columns.
Use AWS Glue crawlers to infer the schemas and available columns.
Use AWS Glue DataBrew for data cleaning and feature engineering.
Step 3:
Select...
Select...
Create an Amazon SageMaker batch transform job for data cleaning and feature engineering.
Store the resulting data back in Amazon S3.
Use Amazon Athena to infer the schemas and available columns.
Use AWS Glue crawlers to infer the schemas and available columns.
Use AWS Glue DataBrew for data cleaning and feature engineering.

Question 3. (Multi Select)

An ML engineer is training a simple neural network model. The ML engineer tracks the performance of the model over time on a validation dataset. The model's performance improves substantially at first and then degrades after a specific number of epochs.

Which solutions will mitigate this problem? (Choose two.)

A: Enable early stopping on the model.
B: Increase dropout in the layers.
C: Increase the number of layers.
D: Increase the number of neurons.
E: Investigate and reduce the sources of model bias.

Answer: A, B

Explanation:
Early stopping halts training once the performance on the validation dataset stops improving. This prevents the model from overfitting, which is likely the cause of performance degradation after a certain number of epochs.
Dropout is a regularization technique that randomly deactivates neurons during training, reducing overfitting by forcing the model to generalize better. Increasing dropout can help mitigate the problem of performance degradation due to overfitting.

Question 4. (HOTSPOT)

A company wants to host an ML model on Amazon SageMaker. An ML engineer is configuring a continuous integration and continuous delivery (CI/CD) pipeline in AWS CodePipeline to deploy the model. The pipeline must run automatically when new training data for the model is uploaded to an Amazon S3 bucket.

Select and order the pipeline's correct steps from the following list. Each step should be selected one time or not at all. (Select and order three.)

• An S3 event notification invokes the pipeline when new data is uploaded.

• S3 Lifecycle rule invokes the pipeline when new data is uploaded.

• SageMaker retrains the model by using the data in the S3 bucket.

• The pipeline deploys the model to a SageMaker endpoint.

• The pipeline deploys the model to SageMaker Model Registry.

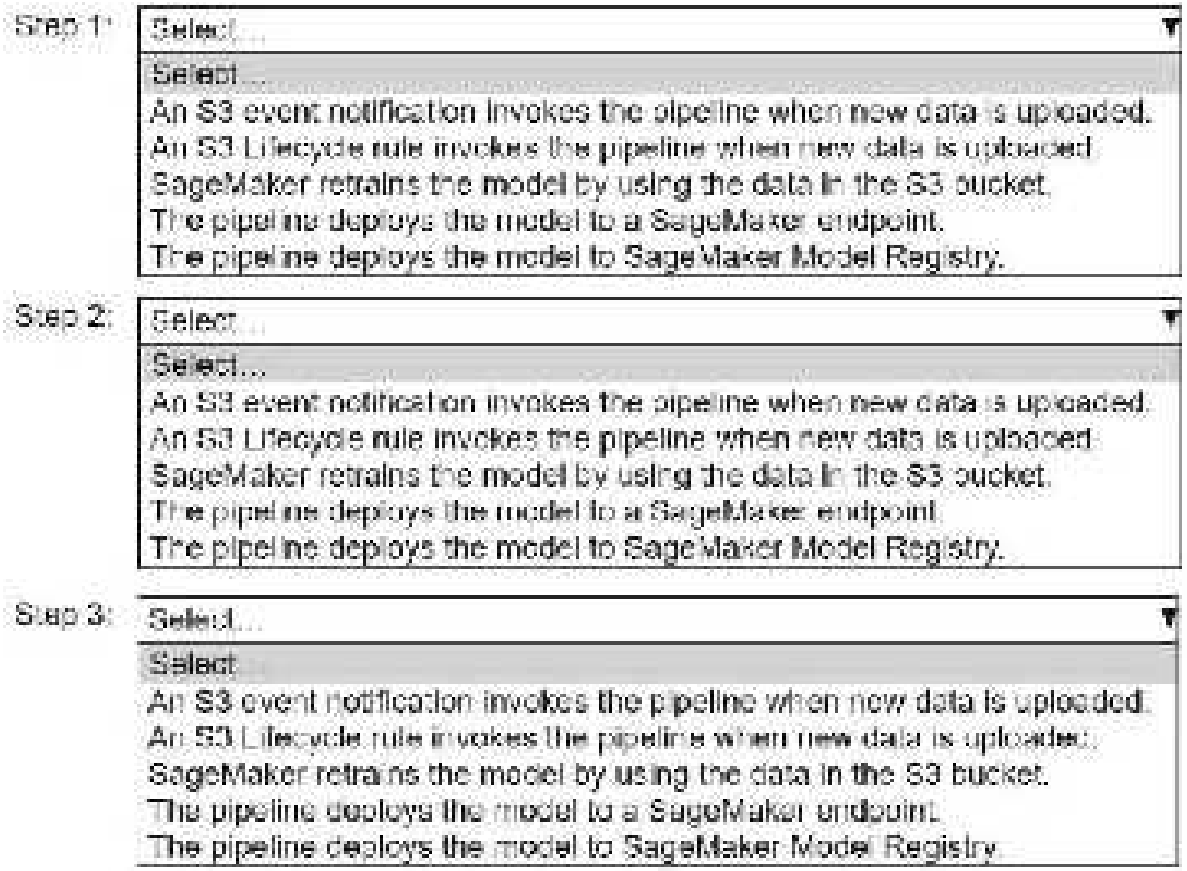
Step 1:
Select...
Select...
An S3 event notification invokes the pipeline when new data is uploaded.
An S3 Lifecycle rule invokes the pipeline when new data is uploaded.
SageMaker retrains the model by using the data in the S3 bucket.
The pipeline deploys the model to a SageMaker endpoint.
The pipeline deploys the model to SageMaker Model Registry.
Step 2:
Select...
Select...
An S3 event notification invokes the pipeline when new data is uploaded.
An S3 Lifecycle rule invokes the pipeline when new data is uploaded.
SageMaker retrains the model by using the data in the S3 bucket.
The pipeline deploys the model to a SageMaker endpoint.
The pipeline deploys the model to SageMaker Model Registry.
Step 3:
Select...
Select...
An S3 event notification invokes the pipeline when new data is uploaded.
An S3 Lifecycle rule invokes the pipeline when new data is uploaded.
SageMaker retrains the model by using the data in the S3 bucket.
The pipeline deploys the model to a SageMaker endpoint.
The pipeline deploys the model to SageMaker Model Registry.

Answer:

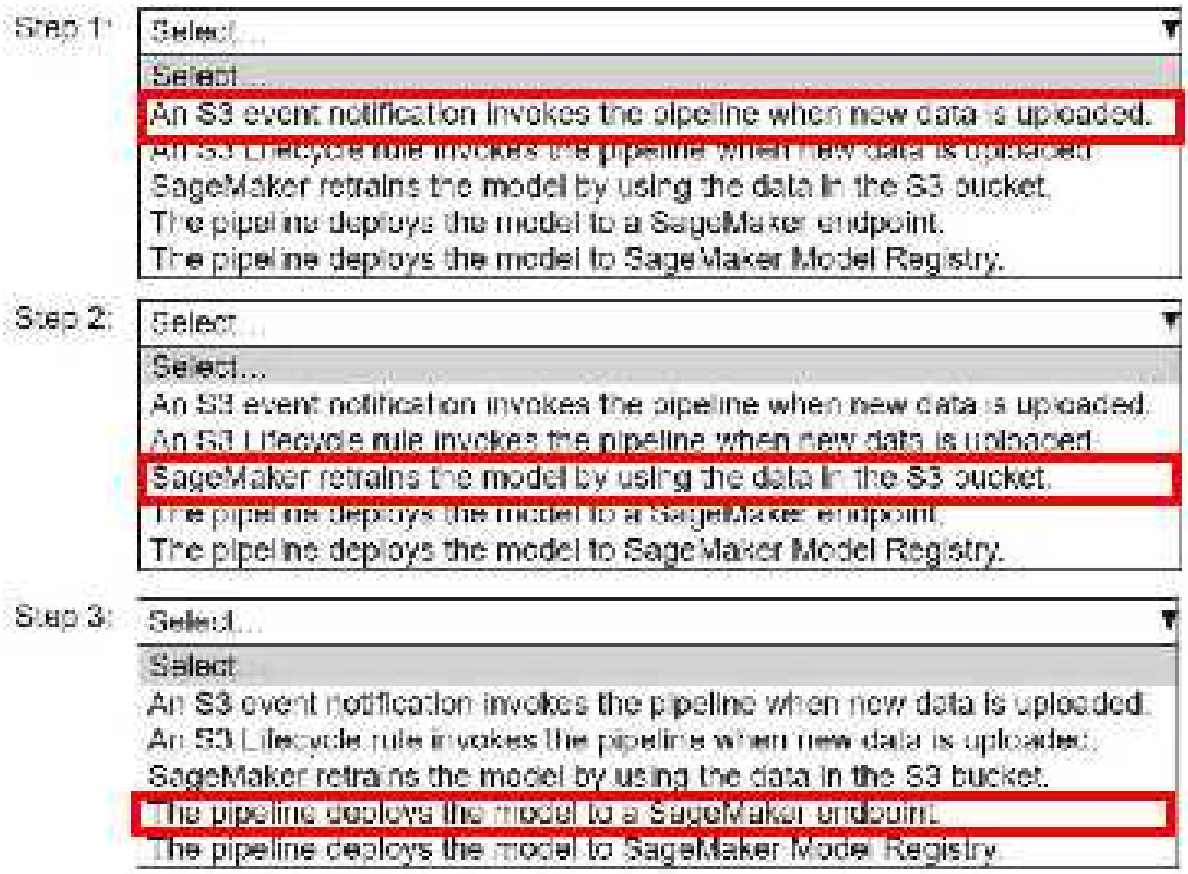
Step 1:
Select...
Select...
An S3 event notification invokes the pipeline when new data is uploaded.
SageMaker retrains the model by using the data in the S3 bucket.
The pipeline deploys the model to a SageMaker endpoint.
The pipeline deploys the model to SageMaker Model Registry.
Step 2:
Select...
Select...
An S3 event notification invokes the pipeline when new data is uploaded.
SageMaker retrains the model by using the data in the S3 bucket.
The pipeline deploys the model to a SageMaker endpoint.
The pipeline deploys the model to SageMaker Model Registry.
Step 3:
Select...
Select...
An S3 event notification invokes the pipeline when new data is uploaded.
SageMaker retrains the model by using the data in the S3 bucket.
The pipeline deploys the model to a SageMaker endpoint.
The pipeline deploys the model to SageMaker Model Registry.

Question 5. (HOTSPOT)

An ML engineer is building a generative AI application on Amazon Bedrock by using large language models (LLMs).

Select the correct generative AI term from the following list for each description. Each term should be selected one time or not at all. (Select three.)

- Embedding
- Retrieval Augmented Generation (RAG)
- Temperature
- Token

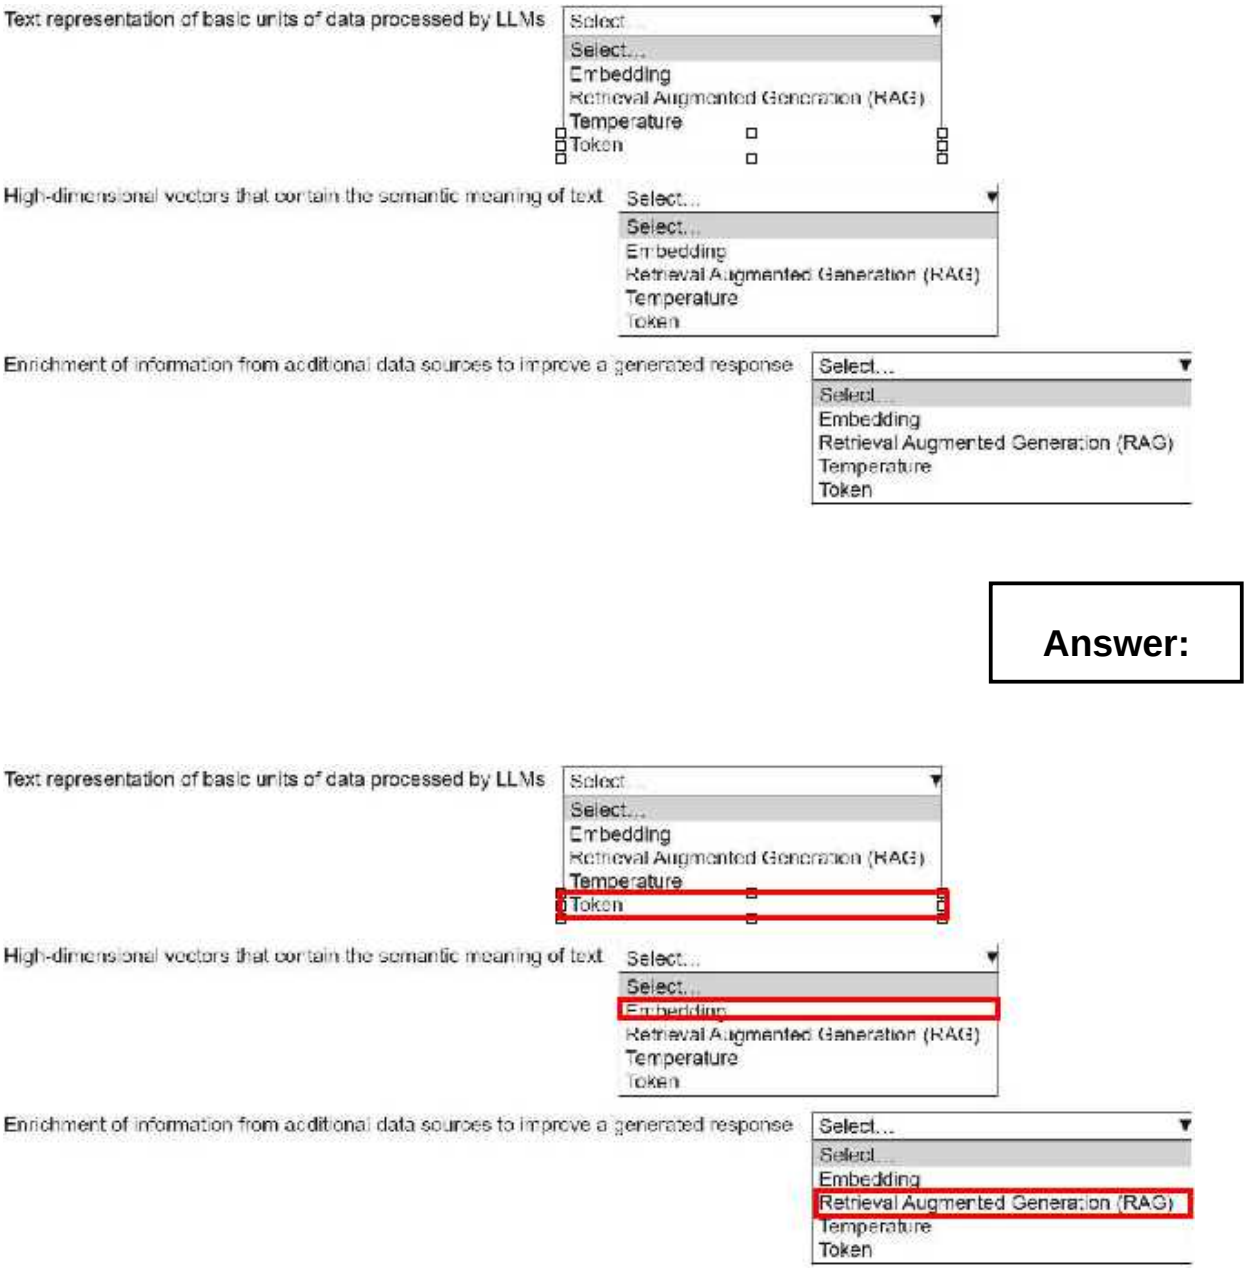
Text representation of basic units of data processed by LLMs
Select...
Select...
Embedding
Retrieval Augmented Generation (RAG)
Temperature
Token
High-dimensional vectors that contain the semantic meaning of text
Select...
Select...
Embedding
Retrieval Augmented Generation (RAG)
Temperature
Token
Enrichment of information from additional data sources to improve a generated response
Select...
Select...
Embedding
Retrieval Augmented Generation (RAG)
Temperature
Token

Answer:

Text representation of basic units of data processed by LLMs
Select...
Select...
Embedding
Retrieval Augmented Generation (RAG)
Temperature
Token
High-dimensional vectors that contain the semantic meaning of text
Select...
Select...
Embedding
Retrieval Augmented Generation (RAG)
Temperature
Token
Enrichment of information from additional data sources to improve a generated response
Select...
Select...
Embedding
Retrieval Augmented Generation (RAG)
Temperature
Token

Need more info? Check the link below:

https://cert4prep.com/exam/mla-c01

Thanks for Being a Valued Cert4Prep User!

Guaranteed Success Pass Every Exam with Cert4Prep.

Save $15 instantly with promo code

LEARN15

Sales: sales@cert4prep.com

Support: support@cert4prep.com